Robotic Graphics: A New Approach to Force Feedback for Virtual Reality

by

William A. McNeely
Boeing Computer Services
P.O. Box 24346
Seattle, WA 98124
wmcneely@espresso.boeing.com

ABSTRACT:

A new conceptual solution is presented for the problem of providing force feedback for virtual reality, concentrating on potential CAD/CAM applications. The essential concept is that force feedback is provided by interactions between the human operator's body and specialized external (as opposed to worn) robots. This is called "robotic graphics" to express the analogy between robots simulating the feel of an object, and graphics displays simulating its appearance. This is illustrated by introducing the derivative concepts of "robotic shape displays" and "roboxels."

The Value of Force Feedback in Virtual Mockup

The development of virtual reality has concentrated on simulating vision, which is reasonable considering how much of the brain is devoted to visual processing. The sense of hearing has received considerably less attention, in part because of the technical difficulty of compensating for individual human hearing characteristics. The tactile sense (stimulating skin receptors only) is relatively straightforward and enjoys ongoing development. The kinesthetic sense ("muscle sense" of contacting massive objects) suffers from severe technical challenges in providing force feedback.

As serious VR applications emerge, the importance of specific types of sensory feedback will be determined by their value in the context of the application. We wish to consider potential CAD/CAM applications of VR. The important senses in conventional CAD/CAM are vision (computer monitor) and tactile and kinesthetic (mouse, keyboard, dial box, etc.). VR promotes inclusiveness and naturalness, offering the hope of interacting in a highly natural way with what appears to be a physical mockup, or what might be called a "virtual mockup."

0-7803-1363-1/93 $3.00 © 1993 IEEE

Analyzing virtual mockup from a technology viewpoint, the conspicuously lagging technology is force feedback. Without it, body and tools pass through solid objects with the greatest of ease! The current approach is to wear mechanical "exoskeletons" that provide passive and active forces as required by the VR scenario. Such devices are complex, bulky, heavy, and potentially dangerous. If an exoskeleton is to simulate immovable objects, then it must be mechanically grounded to the work area, further heightening its undesirable attributes. Researchers at the University of Utah have built a hydraulically actuated exoskeleton arm grounded to a chair [1]. Exos Corporation of Waltham MA has introduced a force-feedback exoskeleton for a single finger grounded to the wrist. A more limited approach is represented by the University of North Carolina's molecule manipulator, in which forces are felt in a gripped manipulator [2].

Faced with this situation, virtual mockup developers will be tempted to avoid supporting force feedback. One can argue, however, that there are no effective sensory substitutes for it in practical applications. For example, consider replacing the spark plugs in a virtual automobile. Even if all contacts between human, spark plug wrench and automobile were immediately indicated using sight, sound and tactile feedback, this normally straightforward task would demand an extremely high level of coordination, concentration and (since limbs tire quickly without support) physical strength. One must normally lean into the engine compartment to reach the least accessible plugs, which becomes outright impossible to simulate. In short, the results of such simulations would seldom be credible.

The above example is drawn from the area of CAD design verification, and specifically, maintenance accessibility. In the area of CAD design origination, the value of force feedback is more subtle, having to do with the need for a good 3-D GUI. For example, a "grab and move" paradigm would certainly be more natural than pointing and clicking, talking, etc. Its naturalness might lead to superior operator efficiency with minimal training, which are major CAD/CAM operational concerns.

Robotic Graphics

Exoskeletons can be seen as an extension of the idea that one simply wears VR devices that provide appropriate sensory feedback. Wearing something, of course, is an excellent way of automatically remaining in close proximity to it. This is certainly practical for small, fixed-shape sensory organs such as eyes and ears, hence, VR headsets. But muscle sense is problematical, because hands, limbs, torso, etc., are relatively large, irregularly shaped and flexible. In that case, the tactic of maintaining continuous body proximity deserves to be weighed against other system considerations such as complexity, bulk, weight and safety.

Myron Krueger has described and implemented an "unencumbered" approach to VR, in which sensory feedback is provided by devices external to the body [3]. While this approach overcomes user objections to wearing devices, it radically sacrifices realism and inclusiveness compared to conventional, encumbered VR. Applying the unencumbered ideal to force feedback, however, suggests that external robots could providethe appropriate forces on a just-in-time basis. Assuming that the user wears a conventional VR headset, such robots would be felt but not seen.

The following ideas elaborate this concept, which we call "robotic graphics" to express the analogy between robots simulating the feel of an object, and graphics displays simulating its appearance.

Robotic Shape Displays (RSD)

The basic idea of what we will call "robotic shape display" (RSD) is illustrated in the following simple VR scenario. A user is seated in front of a virtual desk and wishes to rest an arm on it. The user's arm position is tracked, enabling a running estimation of when and where the user's arm will meet the virtual desktop. A robot is present that can reach any location on the virtual desktop with an end effector consisting of a flat piece of styrofoam ample enough to support an arm. When the system anticipates contact, it orders the robot to bring the styrofoam into just the right position so that the arm comes to rest on it.

To simulate an arm sliding on the desktop, the robot moves the styrofoam so as to keep it constantly positioned underneath the arm. This would not simulate sliding friction, which might be supplied using other means, such as tactile feedback.

More generally, a force-feedback robot could be equipped with a "shape cache," e.g., a turret populated with actual objects that simulate all the different objects that simultaneously appear in a particular VR scenario. As the virtual desktop example illustrates, the cached shapes need not be congruent with the simulated shapes. If the turret is not large enough to accomodate the scene complexity, then it could be dynamically re-populated from a shape repository, with system timeouts as required. Advanced simulations would call for multiple robots.

To simulate a hard immovable object, the robot simply brings an appropriate shape in position and locks its brakes. To simulate a hard moveable object, the robot is augmented with strain sensors that detect user applied forces and, with brakes released, allows the object to move. Viscosity could be simulated in this manner, although the lower the viscosity, the greater the robot system requirements. Object elasticity could be simulated either with active forces, or else with brakes on and using special end effectors containing variable-tension springs.

Critique of RSD

RSD would work best in VR scenarios where the objects are of fixed size and appear repetitively, for example, interacting with a virtual radio's layout of knobs, switches and buttons. It is felt that a large number of useful applications, in manufacturing design, and design verification, are amenable to this approach with existing technology.

By the same token, the most serious theoretical limitation of RSD is that it only supports very simple mathematical shapes (such as flat surfaces) and pre-manufactured objects. To overcome this requires a generalization of RSD that lies well beyond current technology, which is advanced in the discussion of "roboxels."
The shape richness problem might be managed, however, by using simplified shapes to represent

the desired exact ones. For example, exactly displayed objects are enclosed by translucent convex hulls in the visual simulation, and the sides of those hulls are kinesthetically represented by flat-surface end effectors in the manner of the "virtual desktop" example of RSD.

It is interesting to compare such simplified shapes with the CAD technique of creating simplified "access volumes" that surround complex geometry and represent the volume available to a human to perform a given task. RSD simulations of such tasks could be enabled by limiting the access volumes to exactly representable RSD shapes. Alternatively, task simulation could proceed in a two-step process: (1) create a suitable access volume by manipulating RSD shapes, and (2) perform the desired simulation, which is now automatically limited to exactly representable RSD shapes. In general, such simulations would require multi-robot RSD.

RSD simplified shapes could be used for CAD design purposes, by simply grabbing and manipulating them, which would affect the underlying exact object. (As a special case of this, objects could be selected for some CAD operation by simply tapping on them.) Kinesthesia would be based on what happens to the exact objects, not to their RSD simplifications.

RSD requires that user intentions be rapidly and accurately anticipated. Although one can imagine doing this with video cameras or body-mounted position tracking sensors, it nonetheless pushes the limits of current technology. Until reliable and cost-effective capabilities become available, there are interim solutions. The simplest would involve user protocol, e.g., the user first indicates intentions verbally, then waits for contactable object(s) to be visually highlighted before attempting contact.

Miscalculations or system slowness could lead to painful collisions. (Even with an ideally functioning system, the user always has the option of accidentally bumping into things.) The foremost solution is to employ multiple independent safety systems. Robotic safety has been extensively considered in the context of robotic assisted surgery [4]. VR robots could be rated with only enough force for RSD purposes, and the displayed objects could be manufactured from lightweight, collapsible material such as styrofoam. The user could wear protective clothing. Special accident-avoiding software could be used, but it should be considered a secondary line of defense because of software unreliability.

Robot Attendants

Although the drawbacks of RSD could be fairly well managed as described above, the following alternative solution to the problems of tracking and safety is presented. This is presented as a potential refinement of RSD.

The user wears clothing containing rigid attach points. An example of such clothing could be a Velcro-lined arm cuff of the type found in sphygmomanometers, studded with attach points consisting of small ball bearings soldered to circular metallic bases that are sewn into the cuff. Robots could firmly grab and release these attach points, thus providing a way to both track body position and apply passive mechanical resistance.

Consider such a device attached to the forearm. The system is initialized by manually connecting the robot attendant to an attach point. As long as such grip is maintained, it is in principle possible for the robot attendant to track the forearm position with enough accuracy for anticipation purposes, and to restrain the forearm from outpacing the RSD or having an accident.

If the current attach point becomes unsuitable for any reason, then the robot grabs a more suitable attach point (or allows another robot to take over) and releases its grip on the original attach point. If no suitable attach point can be found, then a system limit is reached, and the robot arrests motion.

In principle, such robot attendants could also display the scene-driven forces for body parts that require only gross overall forces, such as torso and limbs. This would place greater requirements on tracking accuracy, clothing design, etc.

A different type of robot attendant would be useful today, in conjunction with the Fake Space Boom (TM) made by Fake Space Labs Inc. of Menlo Park CA. A mobile robot base could ferry the Fake Space Boom (TM) around, allowing an unlimited work area without instrumented rooms. Even with conventional head-mounted displays, a robot attendant could provide tracking over an unlimited work area without instrumented rooms.

Roboxels

The general solution to robotic graphics is to provide cellular robots that dynamically configure themselves into the desired shape and size, lock together and simulate the desired object. We call this a "roboxel," standing for "robotic volume element."

Roboxels are probably several decades away from implementation. It is interesting to note, however, that some of the required infrastructure is currently being pursued in the form of research into micromechanics [5,6] and self-assembling cellular robots [7].

Other aspects of the technical infrastructure for roboxels are waiting to be developed, and it is interesting to speculate about them. It would probably be best to specialize roboxels by function, which would differentiate them by size as well, leading to a hierarchy of robot sizes. By analogy with the body, macroscopic "bone" robots could provide structure. Unlike the body, there is no inherent need for solidity, i.e., there can be hollow pockets in the simulated object. "Nerve" robots could be responsible for communication with the VR computer and with other microrobots. "Skin" microrobotics could be responsible for outer shape and texture. Other roboxels would specialize in strain sensing.

How to supply energy to roboxels? One obvious way would be with batteries, and in fact, researchers at UC-Irvine have produced functioning "microbatteries". Another approach would be to develop "energy" robots containing macroscopic batteries that supply energy to neighboring roboxels over retractable wiring. Still another approach would be for roboxels to absorb electromagnetic energy.

We conclude with some futuristic speculation about other uses for roboxels. Achieving the equivalent of photorealism with roboxels would remove the need for head-mounted visual displays and related paraphernalia. If roboxels became sufficiently cheap to manufacture, they could serve as raw material for manufacturing. Once locked into the desired shape, they could serve as molds, be fused into the finished product, or remain locked together for the useful lifetime of that product and be recycled into a different product.

REFERENCES

[1] P.K. Allen, P. Michelman, K.S. Roberts, "A system for programming and controlling a multi-sensor robotic hand," IEEE Transactions on Systems, Man and Cybernetics, Vol. 20, No. 6, Nov-Dec 1990, pp. 1450-1456.

[2] M. Ouh-young, M. Pique, J. Hughes, N. Srinivasan, F.P. Brooks, Jr., "Using a manipulator for force display in molecular docking," Proceedings of the 1988 International Conference on Robotics and Automation, Vol. 3, pp. 1824-1829, Philadelphia, PA, April 1988.

[3] M. W. Krueger, Artifical Reality 2, Addison-Wesley, 1991.

[4] R.H. Taylor, H.A. Paul, B.D. Mittlestadt, et.al., "An Image Directed Robotic System for Precise Orthopaedic Surgery In Man" Proceedings IEEE Int. Conf. on Engineering in Medicine & Biology, pp. 1928-1930, Philadelphia, PA, Nov. 1990.

[5] Proceedings, IEEE Micro-Electro Mechanical Systems Workshop, Napa Valley, CA. February, 1990.

[6] G. Stix, "Micron Machinations", Scientific American, November 1992, pp. 107-117.

[7] T. Fukuda and Y. Kawauchi, "Cellular Robotic System (CEBOT) as one of the Realization of Self-Organizing Intelligent Universal Manipulator", Proc. IEEE Intl. Conf. on Robotics and Automation, May 1990, pp. 662-667.